REPORT DOCUMENTATION PAGE

Form Approved OMB NO. 0704-0188

1. REPORT DATE (DD-MM-YYYY)	2. REPORT TYPE	3. DATES COVERED (From - To)
30-05-2012	Final Report	1-Sep-2009 - 31-Aug-2011

4. TITLE AND SUBTITLE

Effi□cient Algorithms for Computing Stackelberg Strategies in Security Games

5a. CONTRACT NUMBER
W911NF-09-1-0459

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
611102

6. AUTHORS

Vincent Conitzer, Ronald Parr

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAMES AND ADDRESSES

Duke University
130 Hudson Hall, Box 90271
Duke University
Durham, NC 27705 -

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

U.S. Army Research Office
P.O. Box 12211
Research Triangle Park, NC 27709-2211

10. SPONSOR/MONITOR'S ACRONYM(S)
ARO

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
56698-NS.6

12. DISTRIBUTION AVAILIBILITY STATEMENT

Approved for Public Release; Distribution Unlimited

13. SUPPLEMENTARY NOTES

The views, opinions and/or findings contained in this report are those of the author(s) and should not contrued as an official Department of the Army position, policy or decision, unless so designated by other documentation.

14. ABSTRACT

Game theory provides a framework for modeling a wide range of security and defense problems. This project focuses on Stackelberg strategies, which are optimal when one player can commit to a (possibly randomized) strategy before the other player moves. For example, a defensive unit can commit to a randomized patrolling pattern to deter attacks.

15. SUBJECT TERMS

game theory; algorithms; security

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Vincent Conitzer
UU	UU	UU	UU		19b. TELEPHONE NUMBER: 919-660-6503

Report Title

Effi□cient Algorithms for Computing Stackelberg Strategies in Security Games

ABSTRACT

Game theory provides a framework for modeling a wide range of security and defense problems. This project focuses on Stackelberg strategies, which are optimal when one player can commit to a (possibly randomized) strategy before the other player moves. For example, a defensive unit can commit to a randomized patrolling pattern to deter attacks.

This project explores new approaches for efficiently computing Stackelberg strategies in realistic security domains with exponentially large strategy spaces. Potential impacts of this research include increased ability to compute optimal strategies for security and defense scenarios.

Notable contributions of the project include: (1) New algorithms and complexity results for security games as well as unrestricted games. The algorithms allow us to solve new classes of games efficiently; the complexity results indicate that other methods are needed for richer classes of games. (2) A deeper understanding of the role of commitment and the assumption that the attacker can observe the defender's strategy. These results indicate that, in a sense, Stackelberg strategies are ``safe'' to play even when this assumption does not hold, in some security domains (but not all -- and to address this shortcoming, we also provide a methodology for other security games).

Enter List of papers submitted or published that acknowledge ARO support from the start of the project to the date of this printing. List the papers, including journal references, in the following categories:

(a) Papers published in peer-reviewed journals (N/A for none)

Received	Paper

TOTAL:

Number of Papers published in peer-reviewed journals:

(b) Papers published in non-peer-reviewed journals (N/A for none)

Received	Paper

TOTAL:

Number of Papers published in non peer-reviewed journals:

(c) Presentations

Number of Presentations: 0.00

Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Received	Paper

TOTAL:

Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Peer-Reviewed Conference Proceeding publications (other than abstracts):

Received	Paper	
2012/05/30 1(	1	Vincent Conitzer, Dmytro Korzhyk. Commitment to correlated strategies, Proceedings of the National Conference on Artificial Intelligence (AAAI). 2011/08/07 00:00:00, . : ,
2012/05/30 1(	2	Manish Jain, Dmytro Korzhyk, Ondrej Vanek, Vincent Conitzer, Michal Pechoucek, Milind Tambe. A double oracle algorithm for zero-sum security games on graphs, Proceedings of the Tenth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-11). 2011/05/02 00:00:00, . : ,
2012/05/30 1(	3	Dmytro Korzhyk, Vincent Conitzer, Ronald Parr. Security games with multiple attacker resources, Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI-11). 2011/07/16 00:00:00, . : ,
2012/05/30 1(	4	Dmytro Korzhyk, Vincent Conitzer, Ronald Parr. Solving Stackelberg Games with Uncertain Observability, Proceedings of the Tenth International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-11). 2011/05/02 00:00:00, . : ,

TOTAL: 4

Number of Peer-Reviewed Conference Proceeding publications (other than abstracts):

(d) Manuscripts

Received	Paper	
2012/05/30 1(	5	Dmytro Korzhyk, Zhengyu Yin, Christopher Kiekintveld, Vincent Conitzer, Milind Tambe. Stackelberg vs. Nash in Security Games: An Extended Investigation of Interchangeability, Equivalence, and Uniqueness., Journal of Artificial Intelligence Research (01 2011)

TOTAL: 1

Number of Manuscripts:

Books

Received	Paper

TOTAL:

Patents Submitted

Patents Awarded

Awards

Conitzer:
Thomas Langford Lectureship Award, 2011/2012.
Duke Bass Society of Fellows, 2011-present.
The IJCAI Computers and Thought Award, 2011.
IEEE Intelligent Systems' "AI's 10 to Watch" (Jan./Feb. 2011 issue).
Runner-up for the Best SPC member award, AAMAS 2011.
National Science Foundation CAREER Award, 2010.
Runner-up for the Best SPC member award, AAMAS 2010.
Top 5% of all undergraduate instructors at Duke, Fall 2009.

Graduate Students

NAME	PERCENT_SUPPORTED	Discipline
Dmytro Korzhyk	0.60	
Joshua Letchford	0.27	
FTE Equivalent:	**0.87**	
Total Number:	**2**	

Names of Post Doctorates

NAME	PERCENT_SUPPORTED
FTE Equivalent:	
Total Number:	

Names of Faculty Supported

NAME	PERCENT_SUPPORTED	National Academy Member
Vincent Conitzer	0.09	
Ronald Parr	0.10	
FTE Equivalent:	**0.19**	
Total Number:	**2**	

Names of Under Graduate students supported

NAME	PERCENT_SUPPORTED
FTE Equivalent:	
Total Number:	

Student Metrics

This section only applies to graduating undergraduates supported by this agreement in this reporting period

The number of undergraduates funded by this agreement who graduated during this period:	0.00
The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:.....	0.00
The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:.....	0.00
Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):	0.00
Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:	0.00
The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense	0.00
The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields:	0.00

Names of Personnel receiving masters degrees

NAME

Total Number:

Names of personnel receiving PHDs

NAME

Total Number:

Names of other research staff

NAME PERCENT_SUPPORTED

FTE Equivalent:

Total Number:

Sub Contractors (DD882)

Inventions (DD882)

Scientific Progress

Our JAIR'11 paper (Korzhyk, Yin, Kiekintveld, Conitzer, and Tambe) significantly extends our AAMAS'10 paper. In this paper, we investigate the role of commitment and the assumption that the attacker can observe the defender's strategy; without this assumption, we have a simultaneous-move game and Nash equilibrium would be a more natural solution concept to use. We prove that under a natural restriction on the family of games, defender Stackelberg strategies must also be Nash strategies, and moreover that the Nash equilibria are interchangeable. This interchangeability property means that if one player plays according to one equilibrium and the other player according to another equilibrium, the result is guaranteed to still be an equilibrium. In general games, this is not always true, leading to the dreaded ``equilibrium selection problem'' that a player does not know according to which equilibrium to play---but thanks to the interchangeability property, in these security games we need not worry about choosing the wrong equilibrium, and in particular by the first result we can just choose the Stackelberg strategy. Hence, Stackelberg strategies are robust to changes in the game model that concern commitment and observability. We also ran simulations on games that do not satisfy the properties needed for Stackelberg strategies to also be Nash strategies; the simulations suggest that Stackelberg strategies are still often Nash strategies in these games, except when the attacker can perform complex coordinated attacks in multiple locations.

In an IJCAI'11 paper (Korzhyk, Conitzer, Parr), we further study this problem of an attacker that performs multiple simultaneous attacks. While (as was shown in the JAIR paper) Stackelberg strategies are not usually also Nash strategies in this context, we show that at least the interchangeability property of Nash equilibria is still satisfied, so one still does not need to worry about which equilibrium strategy is the ``right'' one. We also give a polynomial-time algorithm for computing a Nash equilibrium in this context, which initializes the number of defender resources at zero and gradually increases them to the desired number, all the while maintaining an equilibrium of the game. On the other hand, we show that computing a Stackelberg strategy is actually NP-hard. (These results were surprising to us, because, in contrast, in two-player normal-form games, computing a Stackelberg strategy can be done in polynomial time, whereas computing a Nash equilibrium is PPAD-complete and computing an optimal Nash equilibrium is NP-hard.)

Of course, this still does not resolve what to do in such games when one is not sure whether the attacker can observe the mixed strategy (and, hence, whether Stackelberg or Nash is the right model). Our JAIR paper above does propose a game model in which this uncertainty is modeled explicitly, but it does not provide any algorithm for solving these games. In an AAMAS'11 paper (Korzhyk, Conitzer, Parr), we propose an algorithm for solving these games that uses Nash and Stackelberg solvers as subroutines. (The algorithm will work on any game for which such solvers are available.) We show that in simulations a small number of calls to these solvers is sufficient to solve the games.

In another (still unpublished) draft (Letchford, Korzhyk, Conitzer), we study, for various classes of games including security games, how much can be gained by having the ability to commit to a strategy before the other player moves. We find that usually games can be constructed where the gains from commitment are extreme, though when taking an average over many randomly drawn games, the benefits from commitment tend to be much less extreme.

In another AAMAS'11 paper (Jain, Korzhyk, Vanek, Conitzer, Pechoucek, Tambe), we study the ``Mumbai problem'': in response to the 2008 terrorist attacks on Mumbai, the Mumbai police have started to set up checkpoints in

the city; how can we allocate these in a game-theoretically optimal way? We model this (for now) as a zero-sum game between a defender and an attacker on a graph, where the defender chooses edges in the graph to defend and the attacker chooses a target and a path to that target. Crucially, we do allow the targets to have varying values, which makes an earlier exact approach inapplicable; we also show that an existing approximate approach can be arbitrarily suboptimal. We present the RUGGED (Randomization in Urban Graphs by Generating strategies for Enemy and Defender) algorithm, which uses column and constraint generation techniques to incrementally add strategies to the game until convergence to an optimal solution, and show that it scales to the southern part of Mumbai.

In a AAAI'11 paper (Conitzer and Korzhyk), we study the computation of Stackelberg strategies in general normal-form games. We show that there is a close relationship between the standard linear program for computing a correlated equilibrium of a game (a fairly well-known relaxation of the concept of Nash equilibrium), and the linear-programming approach for computing Stackelberg strategies. This suggests a new linear-programming approach for computing Stackelberg strategies, and in our simulations on 50x50 games this new formulation is faster than the standard approach that involves solving multiple LPs. Perhaps more importantly, it gives a way to extend this approach to more than two players -- specifically, to settings with a single leader and an arbitrary number of followers. This generalization to more than two players does require that the leader can send signals to the followers. (Similarly, in a correlated equilibrium, a mediator sends signals to all the players.)

Technology Transfer

Efficient Algorithms for Computing Stackelberg Strategies in Security Games Final Report

Vincent Conitzer and Ronald Parr
Duke University

1 Statement of the problem studied

Game theory provides a framework for modeling a wide range of security and defense problems. This project focuses on Stackelberg strategies, which are optimal when one player can commit to a (possibly randomized) strategy before the other player moves. For example, a defensive unit can commit to a randomized patrolling pattern to deter attacks.

This project explores new approaches for efficiently computing Stackelberg strategies in realistic security domains with exponentially large strategy spaces. Potential impacts of this research include increased ability to compute optimal strategies for security and defense scenarios.

Notable contributions of the project include:

1. New algorithms and complexity results for security games as well as unrestricted games. The algorithms allow us to solve new classes of games efficiently; the complexity results indicate that other methods are needed for richer classes of games.

2. A deeper understanding of the role of commitment and the assumption that the attacker can observe the defender's strategy. These results indicate that, in a sense, Stackelberg strategies are "safe" to play even when this assumption does not hold, in some security domains (but not all – and to address this shortcoming, we also provide a methodology for other security games).

2 Summary of the most important results

Our JAIR'11 paper [5] significantly extends our AAMAS'10 paper. In this paper, we investigate the role of commitment and the assumption that the attacker can observe the defender's strategy; without this assumption, we have a simultaneous-move game and Nash equilibrium would be a more natural solution concept to use. We prove that under a natural restriction on the family

of games, defender Stackelberg strategies must also be Nash strategies, and moreover that the Nash equilibria are interchangeable. This interchangeability property means that if one player plays according to one equilibrium and the other player according to another equilibrium, the result is guaranteed to still be an equilibrium. In general games, this is not always true, leading to the dreaded "equilibrium selection problem" that a player does not know according to which equilibrium to play—but thanks to the interchangeability property, in these security games we need not worry about choosing the wrong equilibrium, and in particular by the first result we can just choose the Stackelberg strategy. Hence, Stackelberg strategies are robust to changes in the game model that concern commitment and observability. We also ran simulations on games that do not satisfy the properties needed for Stackelberg strategies to also be Nash strategies; the simulations suggest that Stackelberg strategies are still often Nash strategies in these games, except when the attacker can perform complex coordinated attacks in multiple locations.

In an IJCAI'11 paper [3], we further study this problem of an attacker that performs multiple simultaneous attacks. While (as was shown in the JAIR paper) Stackelberg strategies are not usually also Nash strategies in this context, we show that at least the interchangeability property of Nash equilibria is still satisfied, so one still does not need to worry about which equilibrium strategy is the "right" one. We also give a polynomial-time algorithm for computing a Nash equilibrium in this context, which initializes the number of defender resources at zero and gradually increases them to the desired number, all the while maintaining an equilibrium of the game. On the other hand, we show that computing a Stackelberg strategy is actually NP-hard. (These results were surprising to us, because, in contrast, in two-player normal-form games, computing a Stackelberg strategy can be done in polynomial time, whereas computing a Nash equilibrium is PPAD-complete and computing an optimal Nash equilibrium is NP-hard.)

Of course, this still does not resolve what to do in such games when one is not sure whether the attacker can observe the mixed strategy (and, hence, whether Stackelberg or Nash is the right model). Our JAIR paper above does propose a game model in which this uncertainty is modeled explicitly, but it does not provide any algorithm for solving these games. In an AAMAS'11 paper [4], we propose an algorithm for solving these games that uses Nash and Stackelberg solvers as subroutines. (The algorithm will work on any game for which such solvers are available.) We show that in simulations a small number of calls to these solvers is sufficient to solve the games.

In another (still unpublished) draft [6], we study, for various classes of games including security games, how much can be gained by having the ability to commit to a strategy before the other player moves. We find that usually games can be constructed where the gains from commitment are extreme, though when taking an average over many randomly drawn games, the benefits from commitment tend to be much less extreme.

In another AAMAS'11 paper [2], we study the "Mumbai problem": in response to the 2008 terrorist attacks on Mumbai, the Mumbai police have started to set up checkpoints in the city; how can we allocate these in a game-

theoretically optimal way? We model this (for now) as a zero-sum game between a defender and an attacker on a graph, where the defender chooses edges in the graph to defend and the attacker chooses a target and a path to that target. Crucially, we do allow the targets to have varying values, which makes an earlier exact approach inapplicable; we also show that an existing approximate approach can be arbitrarily suboptimal. We present the RUGGED (Randomization in Urban Graphs by Generating strategies for Enemy and Defender) algorithm, which uses column and constraint generation techniques to incrementally add strategies to the game until convergence to an optimal solution, and show that it scales to the southern part of Mumbai.

In a AAAI'11 paper [1], we study the computation of Stackelberg strategies in general normal-form games. We show that there is a close relationship between the standard linear program for computing a correlated equilibrium of a game (a fairly well-known relaxation of the concept of Nash equilibrium), and the linear-programming approach for computing Stackelberg strategies. This suggests a new linear-programming approach for computing Stackelberg strategies, and in our simulations on 50x50 games this new formulation is faster than the standard approach that involves solving multiple LPs. Perhaps more importantly, it gives a way to extend this approach to more than two players – specifically, to settings with a single leader and an arbitrary number of followers. This generalization to more than two players does require that the leader can send signals to the followers. (Similarly, in a correlated equilibrium, a mediator sends signals to all the players.)

3 Personnel funded

Besides the PIs (Conitzer and Parr), two Duke Computer Science Ph.D. students have been funded from this award: Dmytro (Dima) Korzhyk and Joshua (Josh) Letchford. Both are currently expected to complete their Ph.D. dissertations on topics closely related to this grant in 2013.

References

[1] Vincent Conitzer and Dmytro Korzhyk. Commitment to correlated strategies. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 632–637, San Francisco, CA, USA, 2011.

[2] Manish Jain, Dmytro Korzhyk, Ondrej Vanek, Vincent Conitzer, Michal Pechoucek, and Milind Tambe. A double oracle algorithm for zero-sum security games on graphs. In *Proceedings of the Tenth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 327–334, Taipei, Taiwan, 2011.

[3] Dmytro Korzhyk, Vincent Conitzer, and Ronald Parr. Security games with multiple attacker resources. In *Proceedings of the Twenty-Second Interna-*

tional Joint Conference on Artificial Intelligence (IJCAI), pages 273–279, Barcelona, Catalonia, Spain, 2011.

[4] Dmytro Korzhyk, Vincent Conitzer, and Ronald Parr. Solving Stackelberg games with uncertain observability. In *Proceedings of the Tenth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1013–1020, Taipei, Taiwan, 2011.

[5] Dmytro Korzhyk, Zhengyu Yin, Christopher Kiekintveld, Vincent Conitzer, and Milind Tambe. Stackelberg vs. Nash in security games: An extended investigation of interchangeability, equivalence, and uniqueness. *Journal of Artificial Intelligence Research*, 41(2):297–327, 2011.

[6] Joshua Letchford, Dmytro Korzhyk, and Vincent Conitzer. On the value of commitment. *Draft*, 2012.